

FILE CODE


2

SECURITY CLASSIFICATION OF THIS PAGE

AD-A216 578

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT DISTRIBUTION UNLIMITED	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S)		5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-LK. 89-1851	
6a. NAME OF PERFORMING ORGANIZATION Department of Psychology University of Massachusetts	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION AFOSR/NL	
6c. ADDRESS (City, State and ZIP Code) Tobin Hall Amherst, MA 01003		7b. ADDRESS (City, State and ZIP Code) Building 410 Bolling AFB, DC 20332	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION AFOSR	8b. OFFICE SYMBOL (If applicable) NL	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER AFOSR-83-0215	
8c. ADDRESS (City, State and ZIP Code) Building 410 Bolling AFB, DC 20332		10. SOURCE OF FUNDING NOS.	
		PROGRAM ELEMENT NO. 61102F	PROJECT NO. 2312
		TASK NO. A1	WORK UNIT NO.
11. TITLE (Include Security Classification) Biological Investigations of Adaptive Networks: Neuronal Control of Conditioned Responding			
12. PERSONAL AUTHOR(S) John W. Moore			
13a. TYPE OF REPORT Forecast Report Final		13b. TIME COVERED FROM 5/1/84 TO 11/1/84	14. DATE OF REPORT (Yr., Mo., Day) 27 Nov 84
		15. PAGE COUNT 14	
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB. GR.	
19. ABSTRACT (Continue on reverse if necessary and identify by block number)			
<div style="text-align: right;"> DTIC ELECTE JAN 5 1990 S B D </div>			
<div style="border: 1px solid black; padding: 5px; display: inline-block;"> DISTRIBUTION STATEMENT A Approved for public release; Distribution Unlimited </div>			
90 01 04 182			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT CLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS <input type="checkbox"/>		21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL		22b. TELEPHONE NUMBER (Include Area Code)	22c. OFFICE SYMBOL



The Commonwealth of Massachusetts
University of Massachusetts
Amherst 01003

DEPARTMENT OF PSYCHOLOGY

November 27, 1984

*Biological Investigation of Adaptive Networks;
 Neuronal Control of Conditioned Responding*

Dr. William O. Berry
 Program Manager
 Life Sciences Directorate
 Air Force Office of Scientific
 Research (AFSC)
 Bolling Air Force Base, D.C. 20332

Re: AFOSR 83-0215
 (Adaptive Networks)

Dear Bill:

This Research Progress and Forecast Report covers the period from May 1, 1984 to the present. There have been a number of significant developments in our research on adaptive neural networks using the classically conditioned nictitating membrane (NM CR) in rabbit as a model system.

As you know, we are placing a greater emphasis on the theoretical aspects of learning than originally proposed. We seek to develop and evaluate real-time mathematical models that have potential applicability to learning in the AI domain. Our primary goal is to assess computational versions of these models against behavioral and physiological data on the NM CR using simulation experiments. Our most recent efforts have focused on three classes of models.

1. Physiologically Constrained Sutton-Barto Model. This algorithm was originally developed by John E. Desmond. With additional work by Neil Berthier, and with advice from Rich Sutton, this model has been rendered into a form that can make predictions about CR topography and the firing pattern of neurons related to the CR. The original version of the model could do this reasonably well for the case of a single CS paired with the US in a forward-delay paradigm. The model has now been generalized so that it can predict CR topography (or simply associative strength) and single-unit physiological data within complex training paradigms that involve two CSs with independent on-and off-times with respect to each other and the US. These paradigms are analogous to spatial and temporal credit-assignment problems in reinforcement learning. Simulations are now performed on the University's VAXEN computer network, the same system used by Andy Barto and his collaborators in the Department of Computer and Information Science (COINS). A representative simulation of a two-CS conditioning experiment is included with this report.

Psychology, neural response, learning, behavior

7

2. Physiologically Constrained Moore-Stickney Model. The Moore-Stickney model is a real-time rendering of an attentional-associative learning model developed originally by N. J. Mackintosh in England. This model has been applied to classical conditioning in complex (multiple-CS) paradigms and to goal-seeking behavior. Its main success to date has been to describe the effects of hippocampal lesions in associative learning tasks such as these. Nestor Schmajuk and I are completing a paper describing simulation experiments with the latest variant of the Moore-Stickney model. As part of his dissertation project, Schmajuk is applying the model to the problem of CR topography and correlated neuronal activity.

3. Real-Time Pearce-Hall Model. An alternate to the Mackintosh type of attentional learning model was proposed by Pearce and Hall in England. Schmajuk and I have been developing real-time computational variants of the basic Pearce-Hall model for application to network learning problems involving the hippocampus. The results of some simulation experiments under this model have been submitted for publication and additional simulations of hippocampal lesion effects are described in the paper mentioned in the preceding paragraph.

Theoretical activity has been intense this past few months because of the month-long visit by Dr. E. J. Kehoe of Australia. Kehoe and his collaborators have generated much of the behavioral data from the NM CR preparation on multiple-CS effects. His visit in September set the occasion for a series of seven seminars, Computational Learning Models and Test Beds with presentations by Barto, Desmond, Kehoe, Moore and Schmajuk with significant contributions from Neil Berthier and Rich Sutton.

Since Kehoe's departure, we have met formally to consider prediction from several models regarding an experiment performed by my colleague, Joe Ayres, with relevance for real-time models. Ayre's experiment concerned the effects of extending a CS in time beyond the reinforcing event. Sutton, Desmond, and Schmajuk discussed the question from the viewpoint of several models, including: the basic Sutton-Barto Model, a variant of the Sutton-Barto Model designated by Sutton as the Adaptative Heuristic Critique with Discount, Desmond's physiologically constrained version of the Sutton-Barto Model, and Moore-Stickney Model of attentional-associative networks. These formal meetings have been extremely valuable because they (a) enhance the groups' understanding of the various models, (b) familiarize colleagues and students with computational approaches to associative learning, and (c) generate ideas for experiments that can provide clear cut discriminations among models. In this regard, Kehoe and I plan to conduct some behavioral experiments relative to these models in parallel in our respective laboratories.

Most of our theoretical work has relied on the experimental literature on behavioral aspects of the NM CR, i.e., on measures of associative strength and on CR topography. Neuronal schemas that represent the various

2

Dr. William O. Berry
November 27, 1984
Page Three

brain systems that might perform a given computational function within a given model are based primarily on data from lesion studies. For example, in the Moore-Stickney Model the hippocampal formation is portrayed as a system that reduced the salience or associability of stimuli depending on the status of the associative network underlying behavior, an assumption based on observed effects of hippocampal lesions. The behavioral/biological foundations of such schemas will improve as relevant physiological evidence begins to surface. It is for this reason that we are accumulating single-unit recordings concurrently with behavioral testing. Success in this domain demands that physiological recordings include neurons within brain structures known to be essential for the NM CR. (Our Annual Technical Report dated May 20, 1984 covers these points more fully). In addition, it is necessary to have computer-assisted statistical techniques that can quantify the relationship between neuronal firing and the fine-grain details of the behavioral CR. John E. Desmond has developed an impressive arsenal of quantitative tools for this purpose. Examples are included with this report. These same quantitative tools can be extended to quantification of the relationship between theory-generated neuronal firing and the activity of actual CR-related neurons.

There has been an important development concerning brain regions essential for the NM CR. Contrary to published reports as recent as last year, it now appears that cerebellar cortex as well as deep nucleus interpositus is essential for the NM CR. This came to light from lesion experiments by Yeo, Hardiman, and Glickstein at University College London, a group with whom I have collaborated for several years. A lesion confined to a portion of cerebellar cortex known to anatomists as the simplex lobe (a.k.a. hemispheric lobule VI, or HVI, for short) profoundly disrupts a previously acquired NM CR while having no effect on the underlying reflex (UR). In short, lesion of HVI, but no other portion of cerebellar cortex, produces the same deficit as does lesions of "downstream" premotor components of the cerebello-rubro-circuit essential for the CR.

I was able to examine relevant histological material and data during a visit to the London laboratory this past June, and we have reproduced the effect of HVI lesions in our laboratory. These lesion results, and other anatomical considerations, suggest that neurons within HVI are crucial for the learning and execution of the NM CR. We have recently begun recording from single neurons of HVI during behavioral testing. We are especially interested in the relationship between the Purkinje cell activity and the CR because Purkinje cells are featured prominently in several mathematical models of cerebellar learning (e.g., Albus, Gilbert, Marr). Photocopies of the experimental set up for recording from Purkinje cells and sample data obtained by Neil Berthier are included with this report.

on For	
A&I	<input checked="" type="checkbox"/>
iced	<input type="checkbox"/>
ation	<input type="checkbox"/>

By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



Dr. William O. Berry
November 27, 1984
Page Four

In order to establish the involvement of Purkinje cells in the NM CR, it is necessary to record from them during behavioral testing. There are a number of questions we intend to pursue:

1. Is there a relationship between Purkinje cell activity and CR topography?
2. What is the nature of this relationship and how well is it described by computational models?
3. Are CR related Purkinje cells limited to HVI, or do they exist in other regions of the cerebellar cortex?
4. Assuming that Purkinje cells are causally implicated in the NM CR, precisely how do they influence other premotor components of the CR, e.g., nucleus interpositus, red nucleus, and supratrigeminal reticular formation?

I can provide more detailed information regarding any aspect of this report immediately on request.

Cordially,



John W. Moore, Ph.D.
Professor of Psychology
(Neuroscience and Behavior)
Associated Professor of
Computer and Information Sciences

JWM/ac1

Enclosures

7

Note: Figs. 2-4 are real data; Figs 5-6 are theory generated

FIGURE 1. Photograph of rabbit in recording apparatus. This animal has been prepared for brain stem recordings with a hollow pedestal through which recording electrodes are introduced into the brain.

FIGURE 2. Representative CRT tracings (100 ms/div) from single conditioning trials showing (upper traces) NM activity and (lower traces) single SR units. A. An on-unit from Animal 32B to a noise CS+. Note increase in firing about 100 ms before the CR. B. Another on-unit in Animal 32B showing increased activity concurrent with the CR to a tone CS-. C. An off-unit in Animal 17 to a tone CS. Note decrease in firing approximately 50 ms before CR. D. The same unit as in C, but on an extinction trial with no CR. E. An on-unit in Animal 26 showing an increase in activity concurrent with the CR to a noise CS+. F. The same unit as in E during a trial with a tone CS- and no CR.

FIGURE 3. Representative computer generated peristimulus-time histograms and averaged NM activity. These are based on software for the Appl. II developed by F. R. Solomon and his colleagues at Williams College. Vertical bars indicate the CS-US interval; The left-hand bar denotes CS onset. The right-hand bar denotes US onset, if it occurs; it is shown for nonreinforced trials to ease comparisons. A. Summary of 21 CS+ trials with CRs for Animal 32B showing a typical on-unit. The dots indicate statistically significant counts in relation to pre-CS activity. The program provides a menu of tests (t or binomial), significance levels (e.g., .05), and null hypothesis rejection regions (1 or 2 tails). B. The same unit as in A, but on 15 CS- trials with no CRs. C. An off-unit in Animal 17 during 16 CS+ trials. D. The same unit as in C on 10 CS- trials.

FIGURE 4. A. Peristimulus-time histogram and averaged NM response over 33 CS+ trials in Animal 02A. B. The same unit as in A over 12 CS- extinction trials with no CRs. Panels A and B are simply reminders to assist understanding Panels C and D. C. Time-correlogram based on the data in Fig. 3A. D. Time-correlogram based on the data in Fig. 3C.

FIGURE 5. Column 1: Second-order conditioning using two 350 ms CSs and no US. CS1 offset is contiguous with CS2 onset. Initial V's: CS2 = 0.9, CS1 = 0. (A). V as a function of time for a trial in which CSs are serially presented. The function is a result of 50 conditioning trials. Hash marks on the abscissa represent, from left to right, CS1 onset, CS2 onset, CS2 offset. (B). V as a function of trials for CS1 and CS2, plotted over 50 trials. (C). Peristimulus time histogram of simulated neuronal activity accumulated over 50 CS1-CS2 presentations. Hash marks on abscissa are the same as in A. Column 2: Serial compound conditioning using a 700 ms CS1, 350 ms CS2, and 30 ms US. CS1 and CS2 coterminate, and their offset is contiguous with US onset. Initial V = 0 for both CSs. (D). V as a function of time for a compound stimulus presentation after 100 trials. Hash marks on abscissa represent, from left to right, CS1 onset, CS2 onset, US onset (and compound offset). (E). V as a function of trials for CS1 and CS2, plotted over 100 trials. (F).

5

Peristimulus time histogram of simulated neuronal activity accumulated over 100 compound stimulus presentations. Abscissa hash marks are identical to those in D.

FIGURE 4. Conditioned inhibition, in which two trial types are presented. For the first type, a reinforced compound, CS+, consists of a simultaneous presentation of CS1 and CS2. The US onset is continuous with the compound offset. The second trial type consists of an unreinforced presentation of CS2 alone. CS1 and CS2 are both 350 ms in duration. US is 30 ms in duration. Initial V's are 0 for both CS's. (A). V as a function of time for both trial types. Hash marks on the abscissa represent, from left to right, the onset and the offset of either CS+ or CS-. (B). V as a function of trials for CS+ and CS-, plotted over 100 trials. (C). Peristimulus time histogram of simulated neuronal activity to CS- presentations. Spike counts are accumulated over 50 trials starting after the initial 100 trials of training. Abscissa hash marks are as indicated in A. (D). Peristimulus time histogram of simulated neuronal activity to CS+ presentations. Spike counts are accumulated as described in C.




FIGURE 1

7

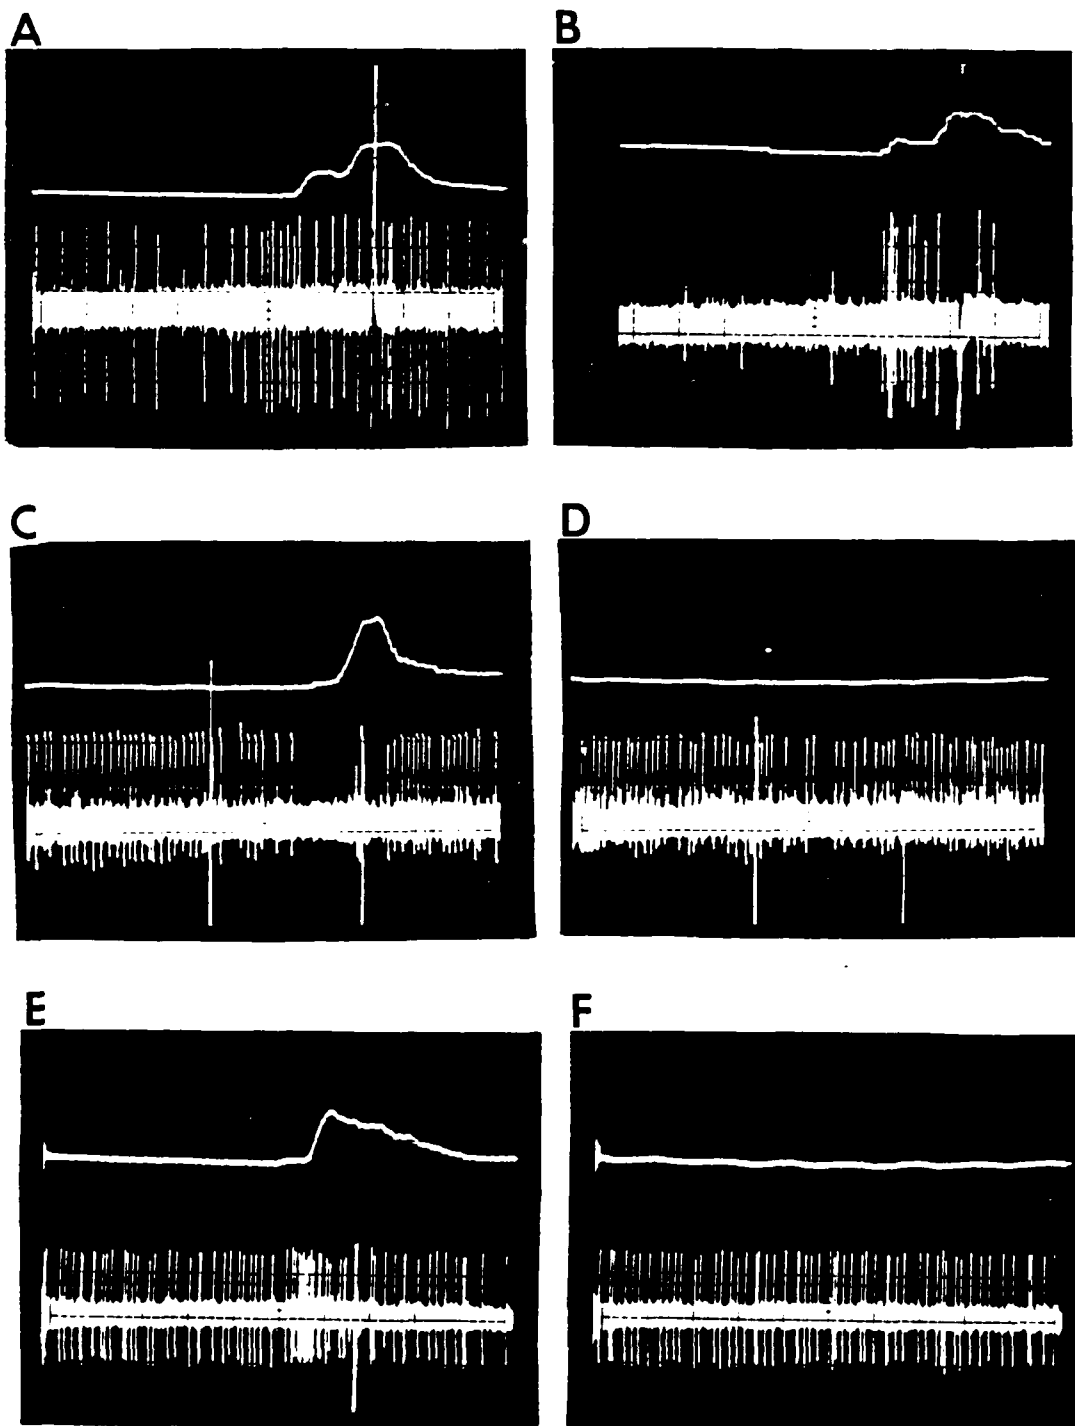



FIGURE 2

8


A



SUB: 02BN+ CR TRIALS=21 CS=350 MS
BIN=10 MS V.CAL.=.2 CNTS, 8 VOLTS

T TEST WAS USED
.05 2-TAILED DF=34
LEAD TIME=50 MS
0 BINS OMITTED
CR ONSET AT 5% MAX AMP
NEURAL ONSET AT SIG. BIN # 9
MEAN SPIKES/TRIAL DEPICTED


B



SUB: 02BT- NO CR TRIALS=15 CS=350 MS
BIN=10 MS V.CAL.=.2 CNTS, 8 VOLTS

T TEST WAS USED
.05 2-TAILED DF=34
LEAD TIME=-370 MS
0 BINS OMITTED
CR ONSET AT 5% MAX AMP
NEURAL ONSET AT SIG. BIN # 1
MEAN SPIKES/TRIAL DEPICTED


C



SUB: 17 TRIALS=16 CS=350 MS
BIN=10 MS V.CAL.=.4 CNTS, 10.66 VOLTS

T TEST WAS USED
.05 2-TAILED DF=34
LEAD TIME=100 MS
0 BINS OMITTED
CR ONSET AT 5% MAX AMP
NEURAL ONSET AT SIG. BIN # 5
MEAN SPIKES/TRIAL DEPICTED

D




SUB: 17++ TRIALS=10 CS=350 MS
BIN=10 MS V.CAL.=.4 CNTS, 10.66 VOLTS

0 BINS OMITTED
MEAN SPIKES/TRIAL DEPICTED

FIGURE 3


A



SUB: 02A CR TRIALS=33 CS=350 MS
BIN=10 MS V.CAL=1.2 CNTS, 8 VOLTS

0 BINS OMITTED
MEAN SPIKES/TRIAL DEFICTED
THIS IS A MERGED FILE


B



SUB: 02A NO CR TRIALS=12 CS=350 MS
BIN=10 MS V.CAL=1.2 CNTS, 8 VOLTS

0 BINS OMITTED
MEAN SPIKES/TRIAL DEFICTED
THIS IS A MERGED FILE


C



SUB: 32EN+ CR CORRELGRAM

ORDINATE: PEARSON R
ABSCISSA: NO. OF BINS THE SPIKE COUNTS WERE SUBTITLED
TO THE RIGHT (+) OR TO THE LEFT (-)

D




SUB: 17 CORRELGRAM

ORDINATE: PEARSON R
ABSCISSA: NO. OF BINS THE SPIKE COUNTS WERE SUBTITLED
TO THE RIGHT (+) OR TO THE LEFT (-)


FIGURE 4

1



change in synaptic weight by trials (scale divider = 1.1)

2



change in synaptic weight by trials (scale divider = 1.1)





FIGURE 5





FIGURE 6

12

UNIVERSITY OF MASSACHUSETTS
AMHERST

MEMORANDUM

FROM John Moore.....DATE September 10, 1984.....
TO Interested Parties.....
SUBJECT Computational Learning Models and Test Beds.....

Jim Kehoe is visiting for one month. He will be giving three lectures on his research into Compound Stimulus Conditioning. His research provides data for developing and evaluating computational models of conditioning, with special emphasis on CR topography.

We are scheduling a series of lunch time lectures beginning Monday, September 10, 1984. Other talks on modelling CR topography and/or related topics are also scheduled.

A tentative schedule appears below: Time: Noon
Place: Room 101, Middlesex House

1. Monday, September 10, 1984
Jim Kehoe, "Serial Compounds"
2. Wednesday, September 12, 1984
Jim Kehoe, "Combination Rules"
3. Monday, September 17, 1984
John Desmond, "CR Topography and Neurophysiological Correlates"
4. Wednesday, September 19, 1984
"Introduction to Attentional Theories", with John Moore
5. Monday, September 24, 1984
Nestor Schmajuk, "Real-Time Attentional Models and CR Topography"
6. Wednesday, September 26, 1984
Andy Barto, "Maybe Layered Networks"
7. Monday, October 1, 1984
Jim Kehoe, "Goodbye and Lots of Luck"

UNIVERSITY OF MASSACHUSETTS
AMHERST

MEMORANDUM

FROM John Moore DATE 20 November 1984
TO Ayres, Barto, Berthier, Desmond, and Schmajuk
SUBJECT Meeting of 14 November 1984

1. Ayres presented the rationale and outcome of an experiment relevant to real-time learning theories of interest to this group. A draft manuscript entitled, "Extending CS Beyond vs. Prior to Reinforcement" is attached.
2. Sutton reported that the data from Ayres' study are basically consistent with the original Sutton-Barto model. However, the derivative model, Adaptive Heuristic Critic with Discount (AHC-D), predicts an asymmetrical outcome such that extension of a CS beyond the US causes a greater reduction of associative value than an equal extension in a forward direction.
3. Schmajuk presented the results of simulation experiments with a version of the Moore-Stickney described in a draft-manuscript entitled, "Variations of CS Effectiveness Revisited: Modified Attentional Models of Classical Conditioning". This model, as does AHC-D, predicts the asymmetry whereby post-US extension of a CS causes greater reductions of associative value than does forward extension of the CS.
4. Desmond presented simulations from the physiologically constrained version of the Sutton-Barto model that he developed with Neil Berthier's help. These simulations also indicated asymmetries as noted above.

The group seemed to agree that:

1. The Ayres' experiment ~~was~~^{is} indeed relevant to real-time models.
2. A simpler experiment with the "blocked" CS would likely yield different anticipated results than those under consideration here.
3. Anticipated outcomes from the point of view of the various models depend critically on parameters. Thus, for example, simulated experiments on CS extensions before and after a CS depend on the duration of the CS, knowledge of the optimal ISI, and/or rates of recruitment and decay of theoretical variables. These considerations should guide choice of designs and parameters of real experiments.